**National Aeronautics and Space Administration**

**Jet Propulsion Laboratory**
California Institute of Technology
Pasadena, California

# Analytics Center Framework for Estimating the Circulation and Climate of the Ocean

**Thomas Huang**
thomas.huang@jpl.nasa.gov

Group Supervisor - Computer Science for Data-Intensive Applications
Strategic Lead - Interactive Data Analytics

Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive, Pasadena, CA 91109-8099, U.S.A.

# What is ECCO?

- **Estimating the Circulation and Climate of the Ocean** (ECCO) is a consortium endeavors to produce the best possible estimates of ocean circulation and its role in climate
- Combining state-of-the-art ocean circulation models with global ocean and sea-ice data in a physically and statistically consistent manner
- ECCO products are being used in studies on ocean variability, biological cycles, coastal physics, water cycle, ocean-cryosphere interactions, and geodesy

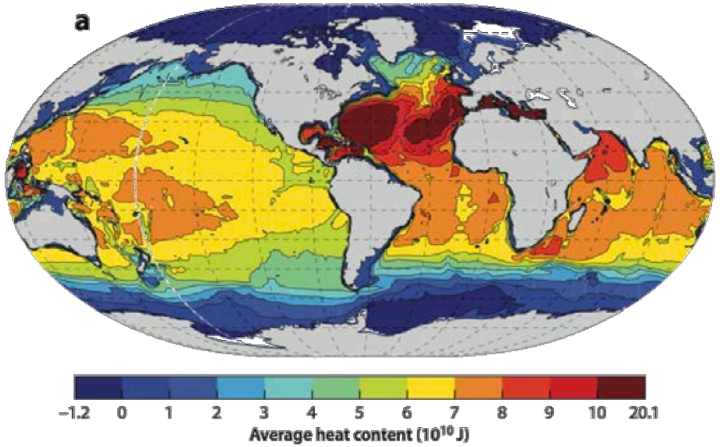


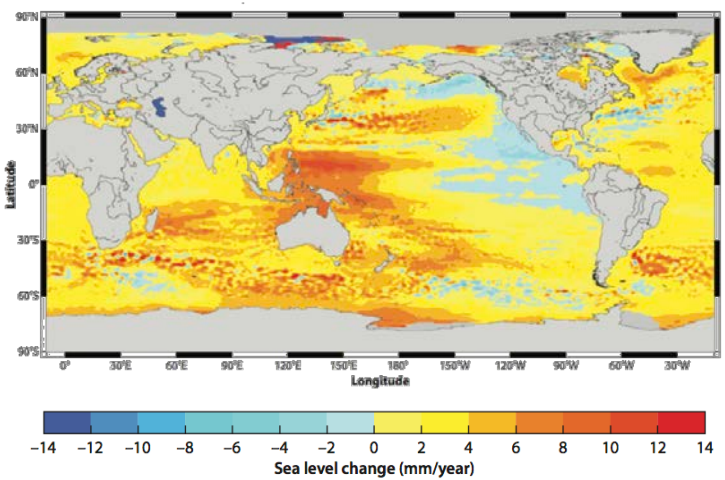

# ECCO Central Production Timeline

- ECCO v4 is the latest release of ocean state estimate
- It is the first adjoint-based, multi-decadal global ocean and sea-ice state estimate

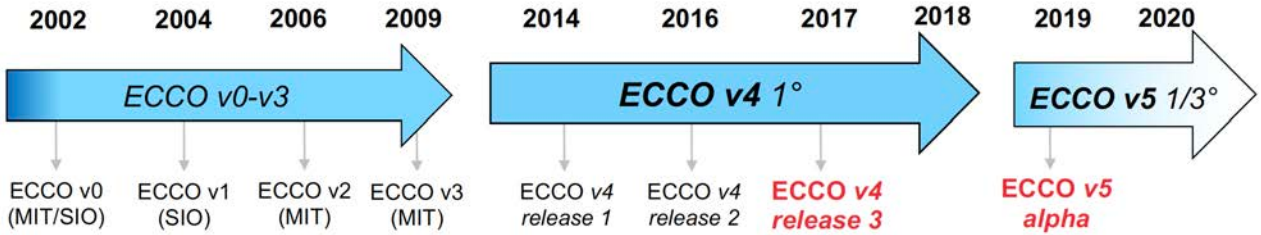


**ECCO v4r3**

- A 1-deg resolution 4D (space and time) **reconstruction** of the **1992 – 2015 global ocean** and **sea-ice state**
- With over 80 variables

- How to visualize and analyze 80 variables?
- How to compare them?
- How to compare ECCO variables with other observational variables to see how they interact?

National Aeronautics and Space Administration

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

# Estimating the Circulation and Climate of the Ocean

NASA ACCESS Program | PI: Patrick Heimbach; Co-Is: Ian Fenty and Thomas Huang

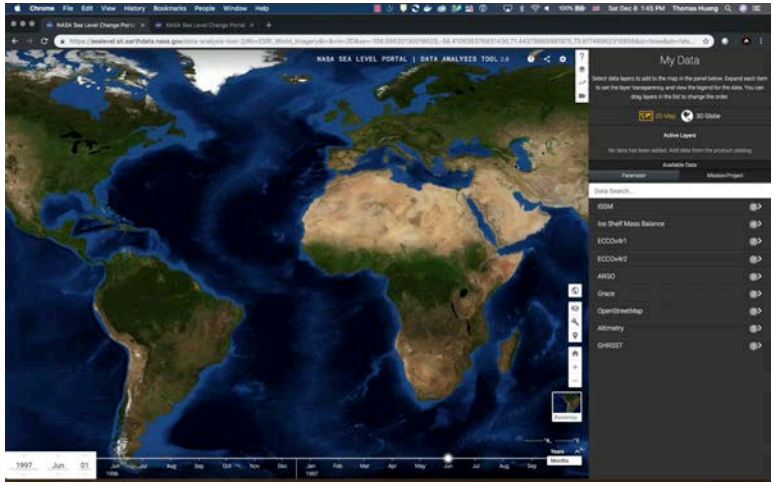


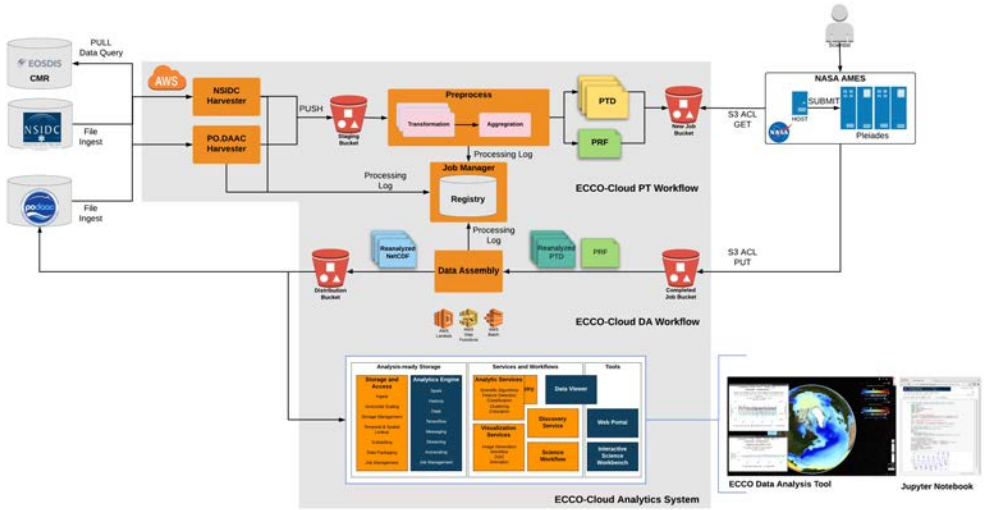


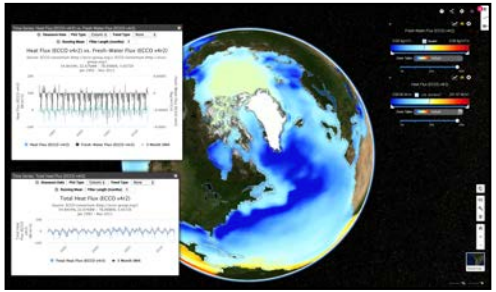

- The **Estimating the Circulation and Climate of the Ocean (ECCO)** global ocean state estimation system (https://ecco.jpl.nasa.gov) is the premier tool for synthesizing NASA's diverse Earth system observations into a complete physical description of Earth's time-evolving full-depth ocean and sea ice system.
- Automate generation of ECCO reanalysis products into CF-compliant netCDF products
- Integrating Amazon Cloud with NASA Ames Pleiades petascale supercomputer
- Establish ECCO Data Analysis Services and web portal for interactive visualization and analysis, and distribution
- Support multi-dimensional data visualization and analysis

# Traditional Method for Analyze Satellite Measurements

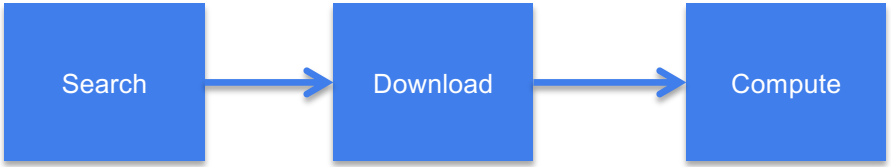


- Depending on the data volume (size and number of files)
- It could take many hours of download – (e.g. 10yr of observational data could yield thousands of files)
- It could take many hours of computation
- It requires expensive local computing resource (CPU + RAM + Storage)
- After result is produced, purge downloaded files

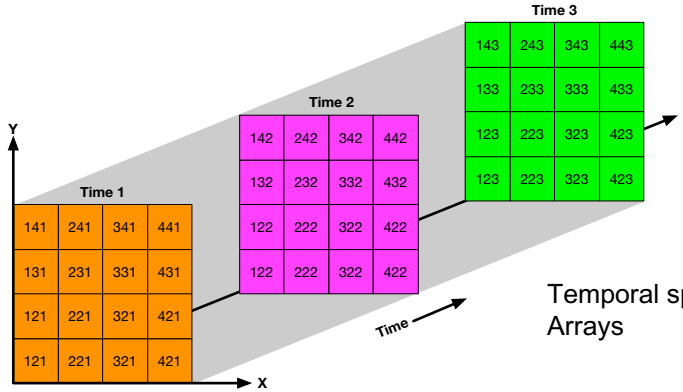


Temporal spatial Arrays

**Observation**

- Traditional methods for data analysis (time-series, distribution, climatology generation) can't scale to handle large volume, high-resolution data. They perform poorly
- Performance suffers when involve large files and/or large collection of files
- A high-performance data analysis solution must be free from file I/O bottleneck

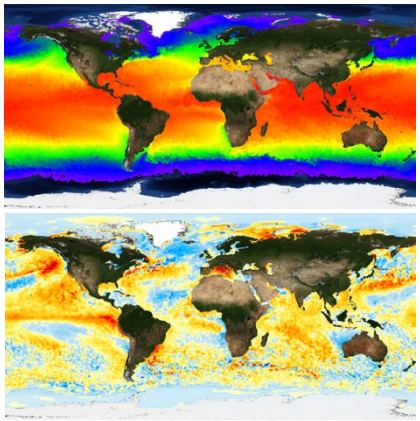

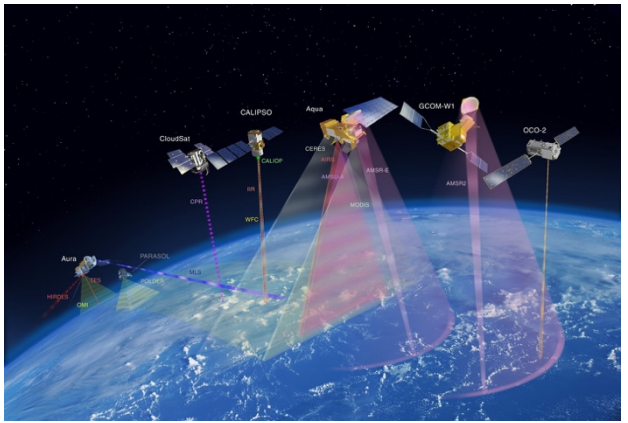

National Aeronautics and Space Administration
Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

# Processors are not Getting Faster

2004: First Pentium 4 processor with 3.0GHz clock speed

2018: Apple's MacBook Pro has clock speed of 2.7GHz

14 years later, not much has gain in raw processing power

**Modern big data architects are required to "think outside of the box". Literally!**

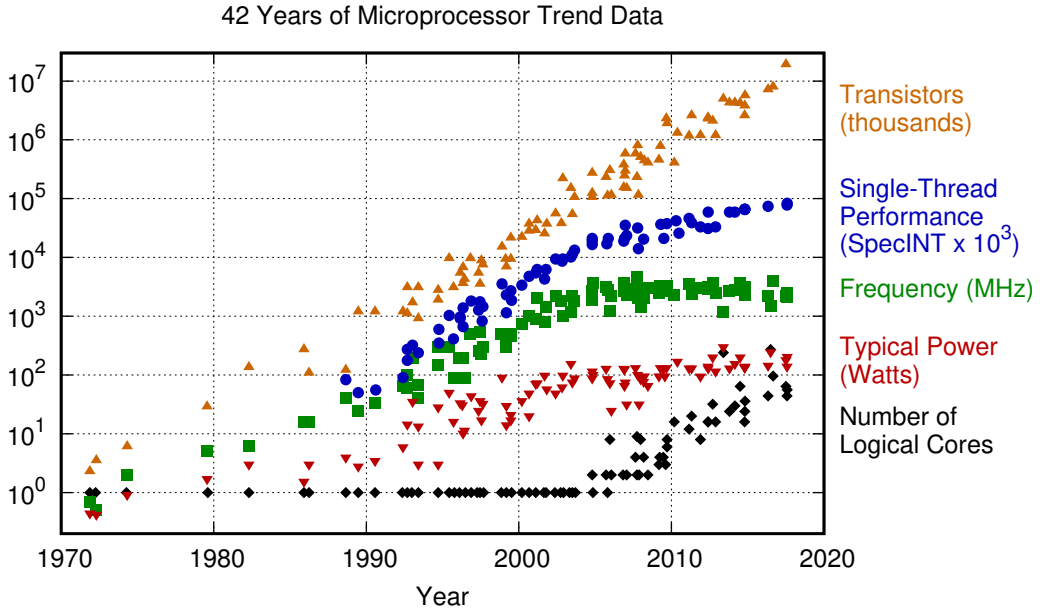


Original data up to the year 2010 collected and plotted by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond, and C. Batten
New plot and data collected for 2010-2017 by K. Rupp

National Aeronautics and Space Administration

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

# Enabling Next Generation of Earth Science Tools and Services

NASA Sea Level Change Portal

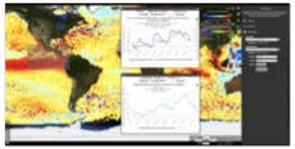

Oceanographic Anomaly Detection

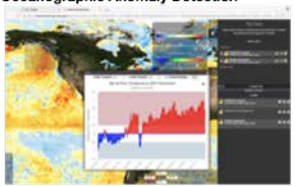

PO.DAAC State Of The Ocean

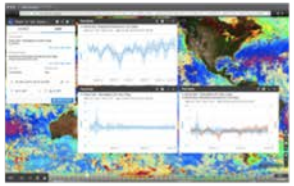

Hydrological Basin Analysis

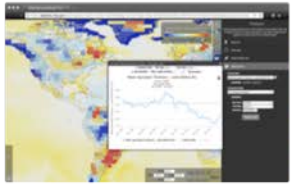

Jupyter Notebook - Interactive Workbench

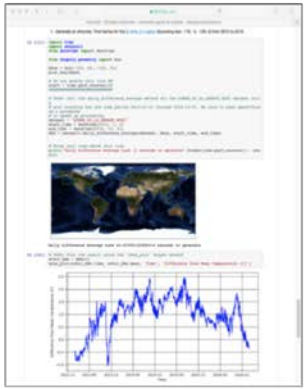

Mobile Analysis

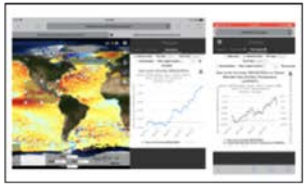

In Situ Data Analysis

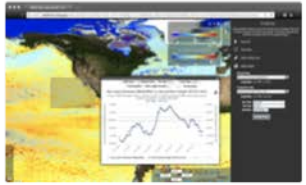

Model Simulations

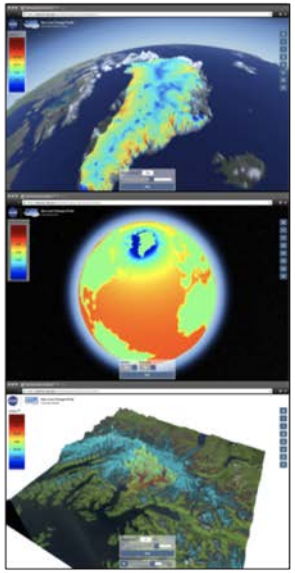

Model - Observation Comparison

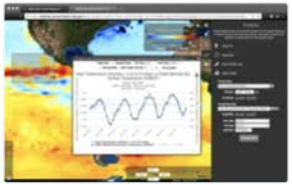

Integrated Search and Discovery


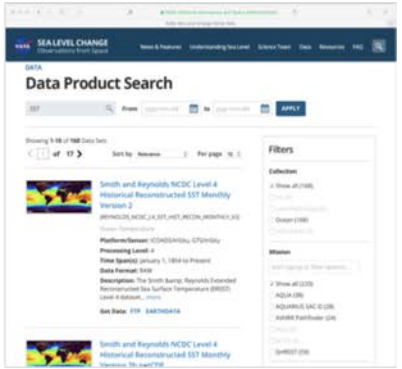

National Aeronautics and Space Administration
Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

# Integrated Science Data Analytics Platform

## Creating SaaS and PaaS for Science Tools and Services

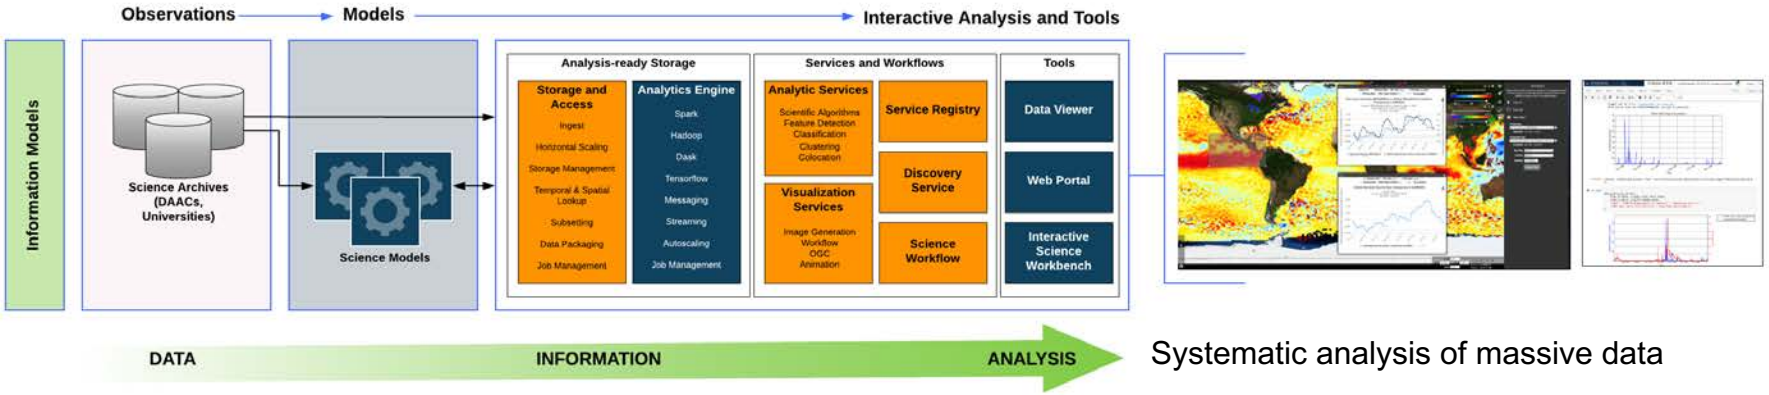


Systematic analysis of massive data

- **Integrated Science Data Analytics Platform**: an analytic center framework to provide an environment for conducting a science investigation
  - Enables the confluence of resources for that investigation
  - Tailored to the individual study area (physical ocean, sea level, etc.)
- Harmonizes data, tools and computational resources to permit the research community to focus on the investigation
- Scale computational and data infrastructures
- Shift towards integrated data analytics
- Algorithms for identifying and extracting interesting features and patterns

National Aeronautics and Space Administration
Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

# Analyze Hurricane Harvey using GPM and SST

## Aug 17, 2017 – Sept. 2, 2017

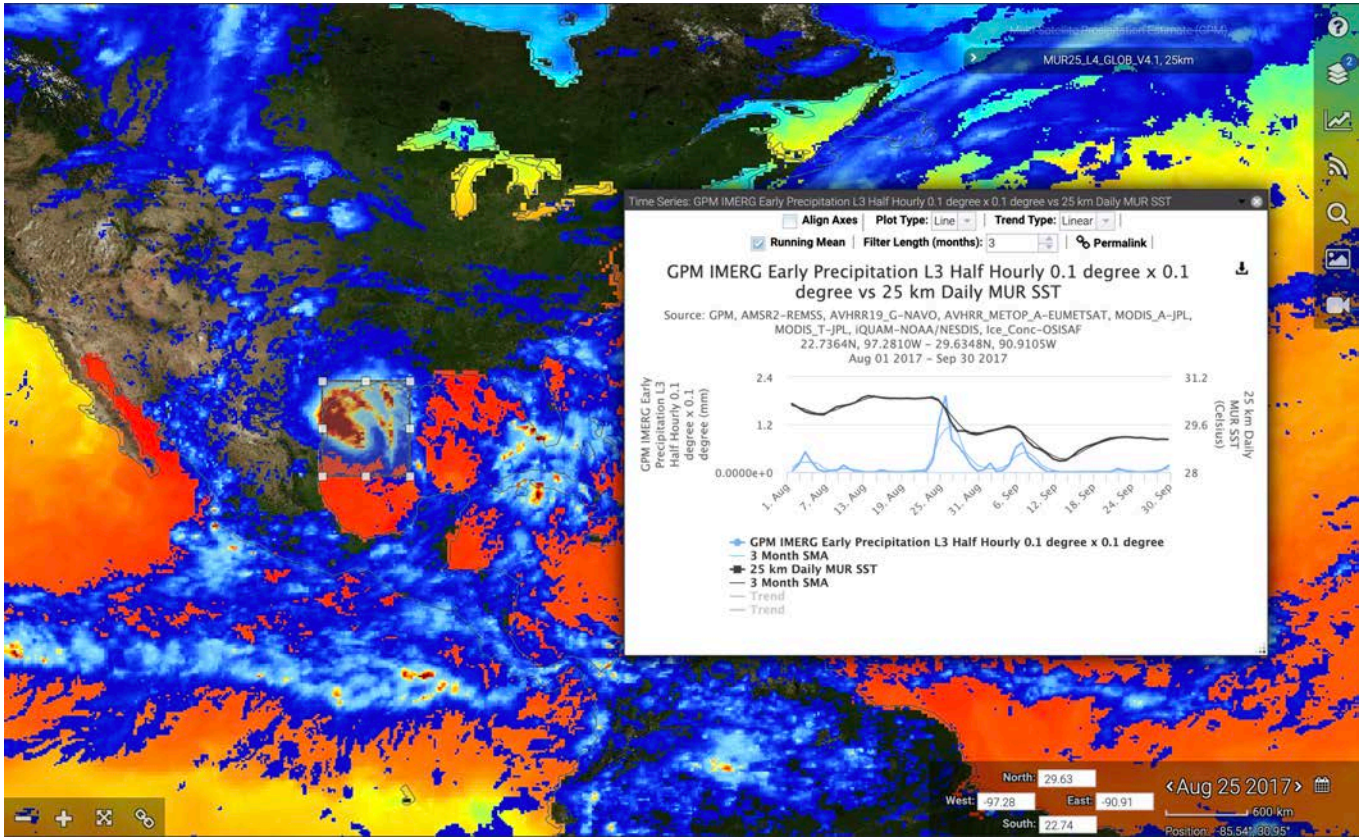

National Aeronautics and Space Administration
Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

# Visualize and Analyze Sea Level

## NASA Sea Level Change Portal - https://sealevel.nasa.gov

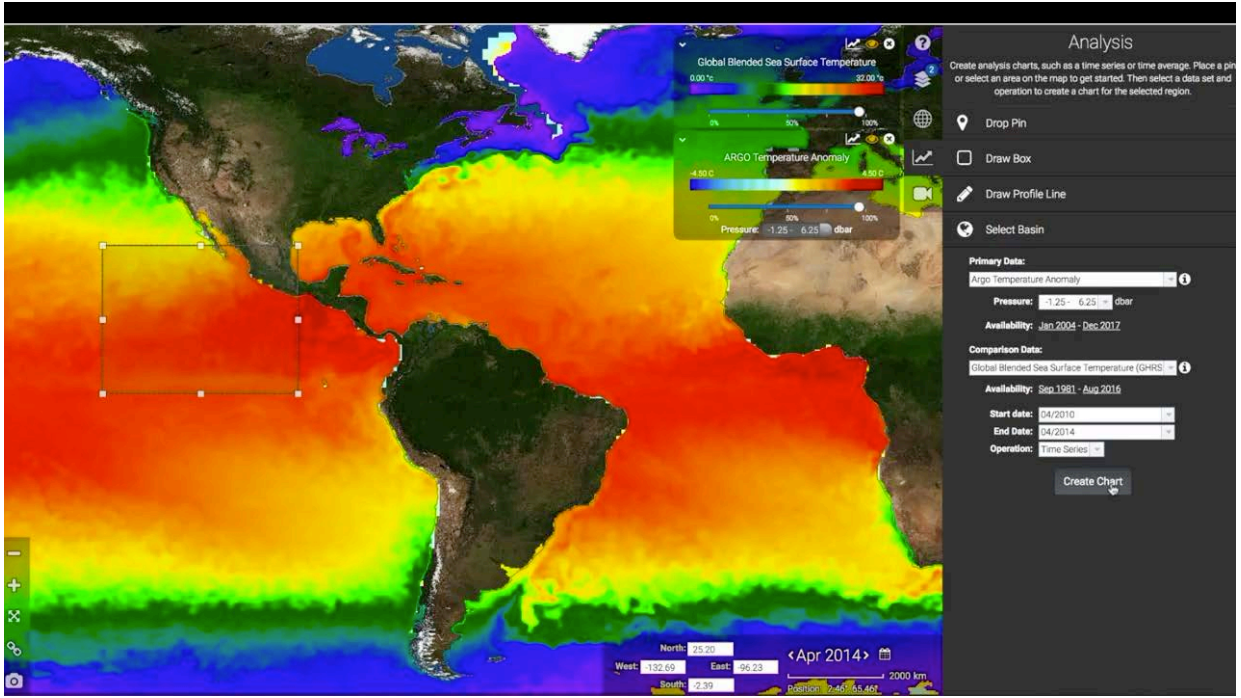


Analyze *in situ* and satellite observations


Analyze Sea Level on mobiles

National Aeronautics and Space Administration
Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

# Hydrology

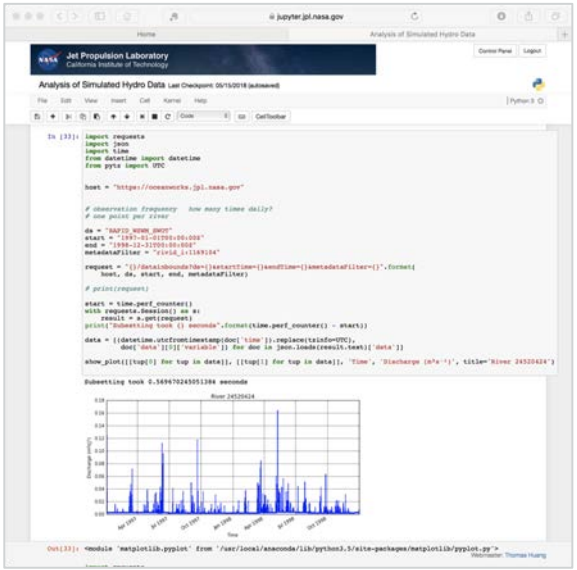


**Retrieval of a single river time series**

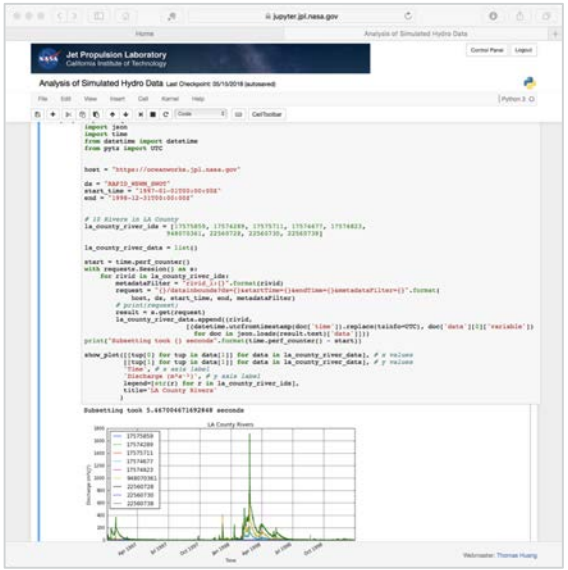


**Retrieval of time series from 9 rivers**

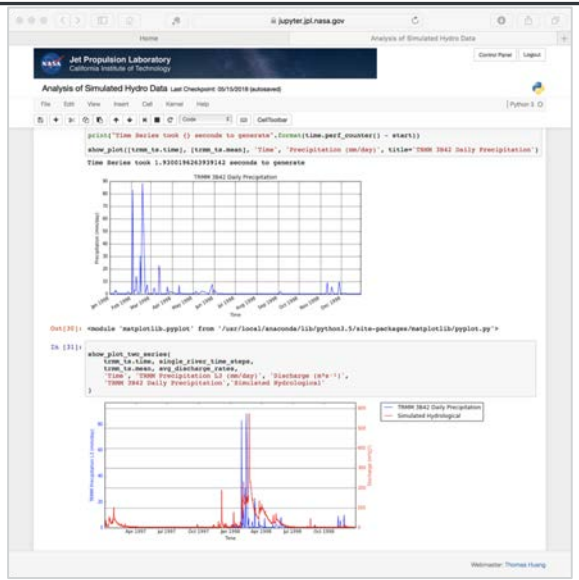


**Time series coordination between TRMM and river**

- Simulated hydrology data in preparation for SWOT hydrology
- **River data**: ~3.6 billion data points. 3-hour sample rate. Consists of measurements from ~600,000 rivers
- **TRMM data**: 17 years, .25deg, 1.5 billion data points
- Sub-second retrieval of river measurements
- On-the-fly computation of time series and generate coordination plot

# Free and Open Source Software (FOSS)

- After more than two years of active development, on October 2017 the **NASA ESOT/AIST OceanWorks** team established Apache Software Foundation and established the **Science Data Analytics Platform (SDAP)** in the **Apache Incubator**
- Technology sharing through Free and Open Source Software (FOSS)
- Why? Further technology evolution that is restricted by projects / missions
- It is more than GitHub
  - Quarterly reporting
  - Reports are open for community review by over 6000 committers
  - SDAP has a group of appointed international mentors
- **SDAP and many of its affiliated projects are now being developed in the open**
  - Support local cluster and cloud computing platform support
  - Fully containerized using Docker and Kubernetes
  - Infrastructure orchestration using Amazon CloudFormation
  - Satellite and model data analysis: time series, correlation map,
  - In situ data analysis and collocation with satellite measurements
  - Fast data subsetting
  - Upload and execute custom parallel analytic algorithms
  - Data services integration architecture
  - OpenSearch and dynamic metadata translation
  - Mining of user interaction and data to enable discovery and recommendations

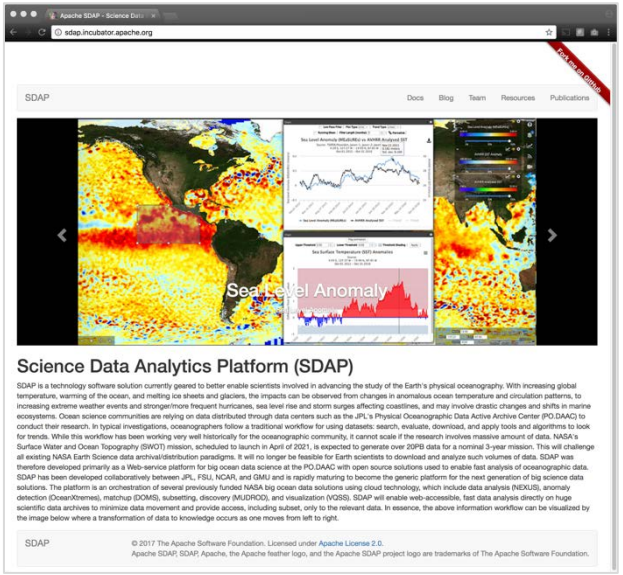


http://sdap.apache.org

APACHE INCUBATOR

**National Aeronautics and Space Administration**

**Jet Propulsion Laboratory**
California Institute of Technology
Pasadena, California

# Other Applications of SDAP

# COVERAGE's Distributed Analytics Center Architecture

- **Committee of Earth Observation Satellites (CEOS) Ocean Variables Enabling Research and Applications for GEO (COVERAGE) Initiative**
- Seeks to provide **improved access** to **multi-agency ocean remote sensing data** that are **better integrated with in-situ and biological observations**, in support of **oceanographic and decision support applications** for societal benefit.
- A community-support open specification with common taxonomies, information model, and API (maybe security)
- Putting value-added services next to the data to eliminate unnecessary data movement
- Avoid data replication. Reduce unnecessary data movement and egress charges
- Public accessible RESTful analytic APIs where computation is next to the data
- Analytic engine infused and managed by the data centers perhaps on the Cloud
- Researchers can perform multi-variable analysis using any web-enabled devices without having to download files

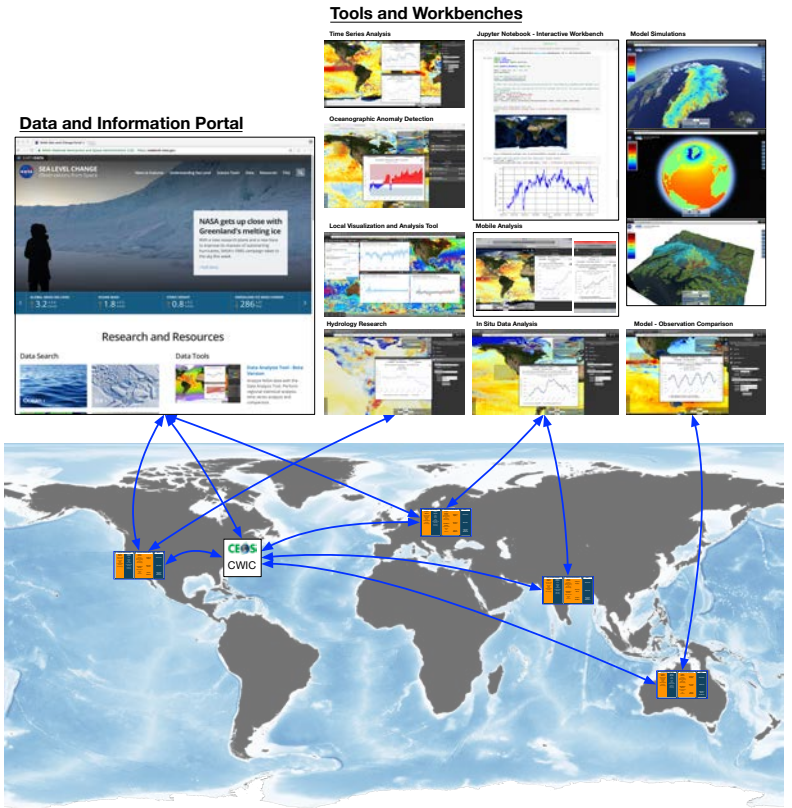

# COVERAGE Phase B

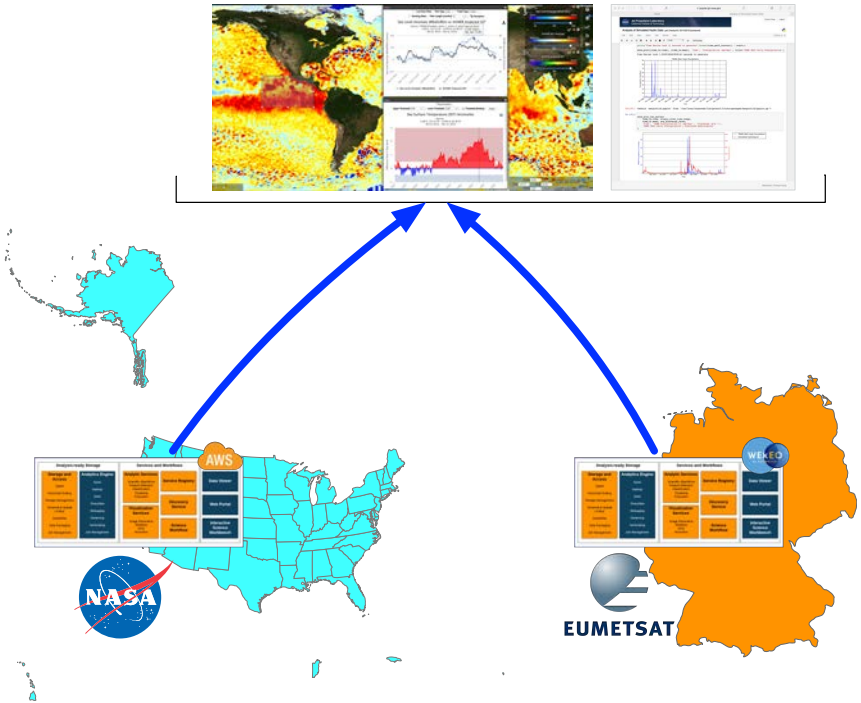


- WEkEO
  - Copernicus Data and Information Access Services (DIAS)
    1. Copernicus Data
    2. Virtual Environment and Tools
    3. User Support
  - Harmonized Data Access for Satellite data and Services
  - Virtualized infrastructure for personal sandboxes
  - Pre-configured tools
- COVERAGE Phase B
  - Establish US Node on Amazon Cloud
  - Establish EU Node on WEkEO at EUMETSAT
  - Establish COVERAGE data portal and analysis tool powered by the COVERAGE Nodes at US and EU

National Aeronautics and Space Administration
Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

# PO.DAAC's SOTO

- NASA's Physical Oceanography Distributed Active Archive Center (PO.DAAC) is an element of the Earth Observing System Data and Information System (EOSDIS)
- PO.DAAC's mission is to preserve NASA's ocean and climate data and make these universally accessible and meaningful
- State of the Ocean (SOTO) is a PO.DAAC's popular visualization tool for the physical oceanography community
- SOTO v5 will be integrated with Apache SDAP and operate on the Amazon Cloud for on-the-fly data analytics

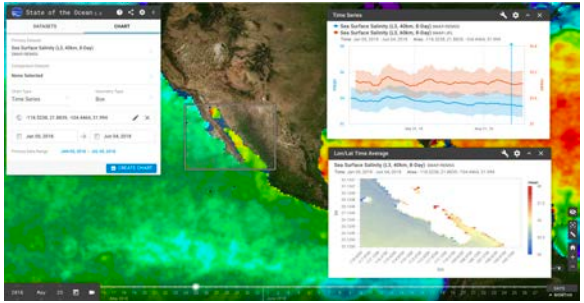

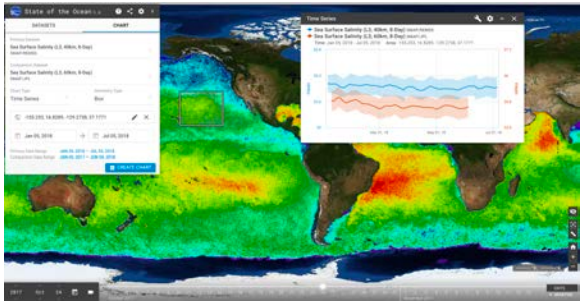

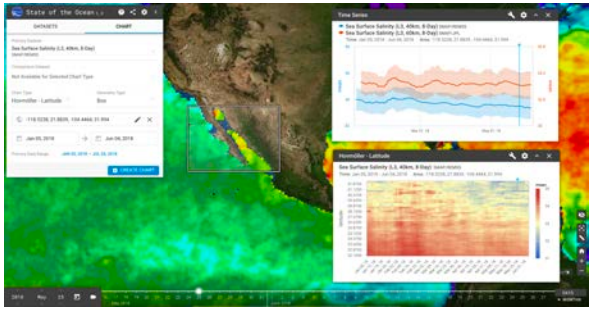

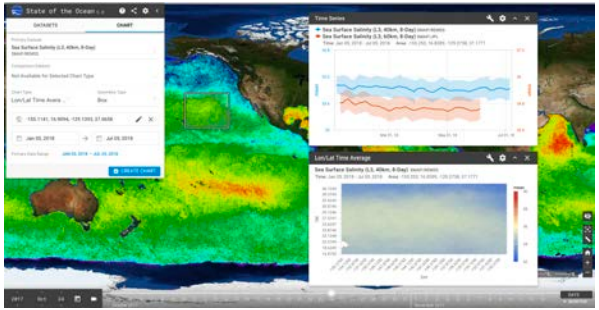

https://podaac.jpl.nasa.gov

National Aeronautics and Space Administration

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

# GRACE-FO Portal and Data Analysis Tool

## GRACE/GRACE-FO Science

- We are developing the new GRACE Follow-On Data Science Portal – https://grace.jpl.nasa.gov
- Goals
  - Common information model
  - Unified data search and access
  - Automated, serverless data processing, analysis and image generation system
  - Integrated with Google Analytics
  - New scientific data analytics capabilities
    - Hydrological basin analysis
    - Regional – country, continent, ocean basin, etc.
    - Multivariate data analysis
  - Deploy on Amazon Web Service with auto-scaling

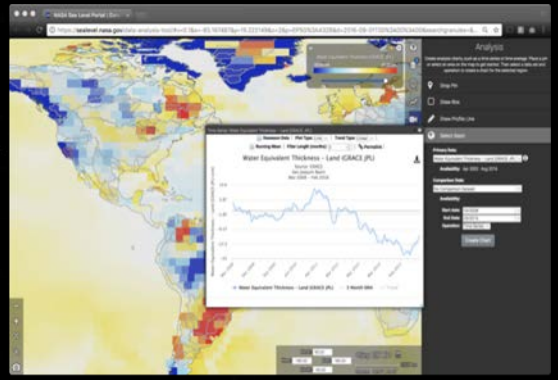

# Building Community-Driven Open Data and Open Source Solutions

- Deliver solutions to establish coherent platform solutions
- Embrace open source software
- Community validation
- Evolve the technology through community contributions
- Share recipes and lessons learned
- Technology demonstrations
- Host webinars, hands-on cloud analytics workshops and hackathons


2019 EGU – NASA Hyperwall

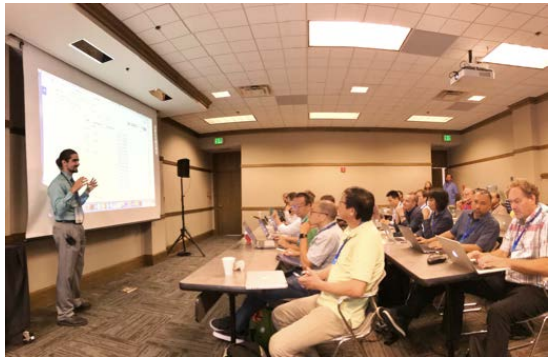

Big Data Analytics and Cloud Computing Workshop, 2017 ESIP Summer Meeting, Bloomington, IN


2019 JPL Data Science Showcase

# Partner with NASA and non-NASA Projects - Deliver to Production

- **The gap between visionary to pragmatists is significant**. – Geoffrey Moore
- Become an expert in the production environment and devote resources in automations
- Give project engineering team early access to the PaaS
- Deliver all technical documents and work with project system engineering
- Provide project-focused trainings

NASA's Sea Level Change Team

CEOS SIT Technical Workshop

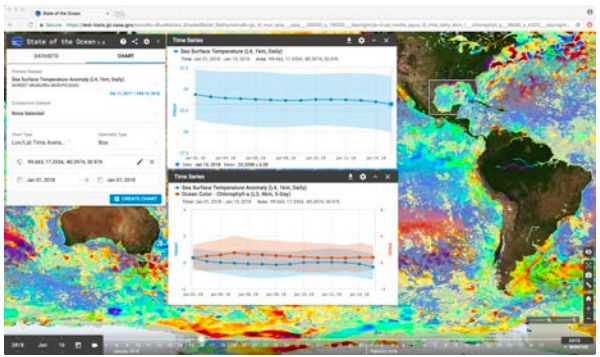

NASA's Physical Oceanography Distributed Active Archive Center (PO.DAAC)

# OceanObs'19

- Our new Community White Paper: "An Integrated Data Analytics Platform"
- https://www.frontiersin.org/articles/10.3389/fmars.2019.00354/full?&utm_source=Email_to_authors_&utm_medium=Email&utm_content=T1_11.5e1_author&utm_campaign=Email_publication&field=&journalName=Frontiers_in_Marine_Science&id=433796
- We are invited to discuss about "Big Ocean Science Analytics using Apache Science Data Analytics Platform" at OceanObs 2019


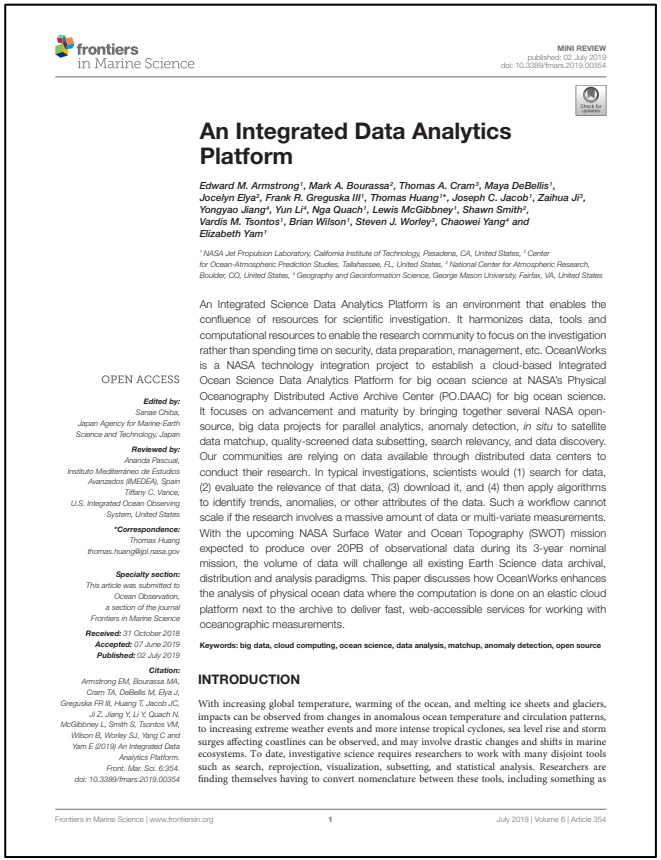

frontiers in Marine Science

MINI REVIEW
published: 02 July 2019
doi: 10.3389/fmars.2019.00354

## An Integrated Data Analytics Platform

*Edward M. Armstrong[1], Mark A. Bourassa[2], Thomas A. Cram[3], Maya DeBellis[1], Jocelyn Elya[2], Frank R. Greguska III[1], Thomas Huang[1]*, Joseph C. Jacob[1], Zaihua Ji[2], Yongyao Jiang[4], Yun Li[4], Nga Quach[1], Lewis McGibbney[1], Shawn Smith[2], Vardis M. Tsontos[1], Brian Wilson[1], Steven J. Worley[3], Chaowei Yang[4] and Elizabeth Yam[1]*

[1] NASA Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA, United States, [2] Center for Ocean-Atmospheric Prediction Studies, Tallahassee, FL, United States, [3] National Center for Atmospheric Research, Boulder, CO, United States, [4] Geography and Geoinformation Science, George Mason University, Fairfax, VA, United States

An Integrated Science Data Analytics Platform is an environment that enables the confluence of resources for scientific investigation. It harmonizes data, tools and computational resources to enable the research community to focus on the investigation rather than spending time on security, data preparation, management, etc. OceanWorks is a NASA technology integration project to establish a cloud-based Integrated Ocean Science Data Analytics Platform for big ocean science at NASA's Physical Oceanography Distributed Active Archive Center (PO.DAAC) for big ocean science. It focuses on advancement and maturity by bringing together several NASA open-source, big data projects for parallel analytics, anomaly detection, *in situ* to satellite data matchup, quality-screened data subsetting, search relevancy, and data discovery. Our communities are relying on data available through distributed data centers to conduct their research. In typical investigations, scientists would (1) search for data, (2) evaluate the relevance of that data, (3) download it, and (4) then apply algorithms to identify trends, anomalies, or other attributes of the data. Such a workflow cannot scale if the research involves a massive amount of data or multi-variate measurements. With the upcoming NASA Surface Water and Ocean Topography (SWOT) mission expected to produce over 20PB of observational data during its 3-year nominal mission, the volume of data will challenge all existing Earth Science data archival, distribution and analysis paradigms. This paper discusses how OceanWorks enhances the analysis of physical ocean data where the computation is done on an elastic cloud platform next to the archive to deliver fast, web-accessible services for working with oceanographic measurements.

**Keywords: big data, cloud computing, ocean science, data analysis, matchup, anomaly detection, open source**

OPEN ACCESS

***Edited by:***
*Sanae Chiba,*
*Japan Agency for Marine-Earth Science and Technology, Japan*

***Reviewed by:***
*Ananda Pascual,*
*Instituto Mediterráneo de Estudios Avanzados (IMEDEA), Spain*
*Tiffany C. Vance,*
*U.S. Integrated Ocean Observing System, United States*

****Correspondence:***
*Thomas Huang*
*thomas.huang@jpl.nasa.gov*

***Specialty section:***
*This article was submitted to Ocean Observation, a section of the journal Frontiers in Marine Science*

***Received:*** *31 October 2018*
***Accepted:*** *07 June 2019*
***Published:*** *02 July 2019*

***Citation:***
*Armstrong EM, Bourassa MA, Cram TA, DeBellis M, Elya J, Greguska FR III, Huang T, Jacob JC, Ji Z, Jiang Y, Li Y, Quach N, McGibbney L, Smith S, Tsontos VM, Wilson B, Worley SJ, Yang C and Yam E (2019) An Integrated Data Analytics Platform. Front. Mar. Sci. 6:354. doi: 10.3389/fmars.2019.00354*

## INTRODUCTION

With increasing global temperature, warming of the ocean, and melting ice sheets and glaciers, impacts can be observed from changes in anomalous ocean temperature and circulation patterns, to increasing extreme weather events and more intense tropical cyclones, sea level rise and storm surges affecting coastlines can be observed, and may involve drastic changes and shifts in marine ecosystems. To date, investigative science requires researchers to work with many disjoint tools such as search, reprojection, visualization, subsetting, and statistical analysis. Researchers are finding themselves having to convert nomenclature between these tools, including something as

# ApacheCon North America

- Invited to discuss our Apache Science Data Analytics Platform (SDAP) project at the ApacheCon North America


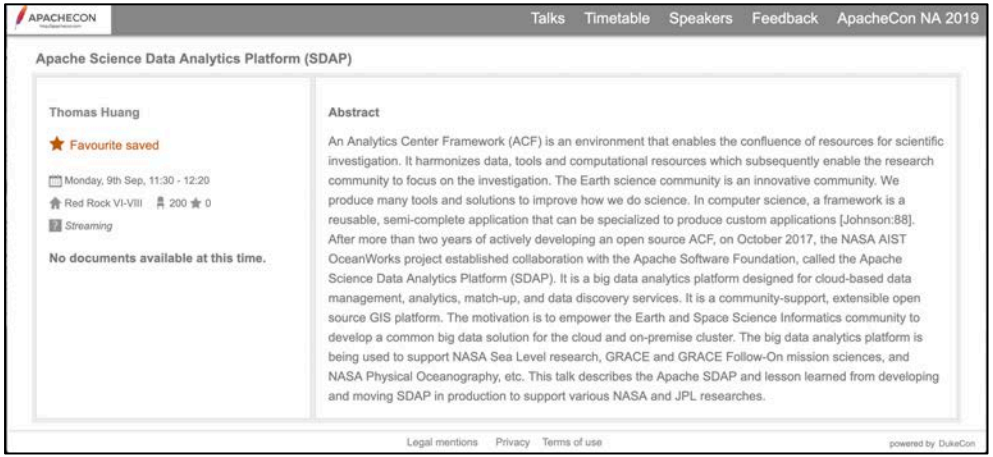

# Acknowledgement

Ed Armstrong/JPL
Jason Barnett/LARC
Andrew Bingham/JPL
Carmen Boening/JPL
Mark Bourassa/FSU
Mike Chin/JPL
Marge Cole/NASA
Tom Cram/NCAR
Dan Crichton/JPL
Maya DeBellis/JPL
Rich Doyle/JPL
Jocelyn Elya/FSU
Ian Fenty/JPL
Kevin Gill/JPL
Frank Greguska/JPL
Patrick Heimbach/UT Austin
Ben Holt/JPL
Thomas Huang/JPL
Joe Jacob/JPL
Zaihua Ji/NCAR
Yongyao Jiang/GMU
Felix Landerer/JPL
Yun Li/GMU
Eric Lindstrom/NASA
Mike Little/NASA
Chris Lynnes/NASA
Lewis McGibbney/JPL
David Moroni/JPL
Kevin Murphy/NASA
Charles Norton/JPL
Jean-Francois Piolle/IFREMER
Nga Quach/JPL
Brandi Quam/NASA
Shawn Smith/FSU
Ben Smith/JPL
Adam Stallard/FSU
Rob Toaz/JPL
Vardis Tsontos/JPL
Suresh Vannan/JPL
Jorge Vazquez/JPL
Ou Wang/JPL
Brian Wilson/JPL
Steve Worley/NCAR
Elizabeth Yam/JPL
Phil Yang/GMU

# In Summary

- **You've got to think about big things while you're doing small things, so that all the small things go in the right direction** – Alvin Toffler
- Focus on end-to-end data and computation architecture, and the total cost of ownership
- JPL Strategy is to drive Data Science into the fabric of JPL by
  - Launching cross-institution pilots
  - Building a trained workforce
  - Linking to the mission-science data lifecycle
- Invest in Interactive Analytics that simplifies the integration of *multiple* Earth observing remote sensing instruments; comparison against models

# Presentation | Evaluation | Collaboration




Presentation | Evaluation | Collaboration

**Thomas Huang**

thomas.huang@jpl.nasa.gov

Jet Propulsion Laboratory

California Institute of Technology